

WHAT IS CLAIMED IS:

1. A method of monitoring an execution of a query, wherein the query execution comprises at least one execution thread, the method comprising the steps of:
  - 5 storing an execution trace record for each execution thread, the execution trace record having execution trace information including a thread identifier (ID) and a time stamp;
  - retrieving the execution trace information; and
  - presenting the retrieved execution trace information to a user after execution
- 10 of the query.
2. The method of claim 1, further comprising the step of executing the query, and wherein the step of storing an execution trace record for each execution thread is performed while executing the query.
- 15 3. The method of claim 1, wherein the query execution comprises a plurality of execution threads, and the step of presenting the retrieved execution trace information to the user comprises the step of synchronizing the execution trace records according to the thread ID and the time stamp.
- 20 4. The method of claim 1, wherein the time stamp is an absolute time stamp.
5. The method of claim 1, wherein the time stamp is a logical time stamp.
- 25 6. The method of claim 1, wherein the step of storing an execution trace record for each execution thread in at least one execution log file comprises the step of splitting the execution log file into a plurality of partitions.

7. The method of claim 6, wherein the execution log file is split into the plurality of partitions based upon an amount of execution trace information and an amount of storage space.

8. The method of claim 6, further comprising the steps of reusing the partitions when the execution log file exceeds a partition size.

9. The method of claim 1, wherein the step of presenting the retrieved execution trace information to a user comprises the steps of:  
accepting a presentation command; and  
arranging the retrieved execution trace records according to the thread ID, the time stamp and the playback command.

10. The method of claim 9, wherein the presentation command is selected from the group comprising a play command, a stop command, a reverse play command, a fast play command, and a fast reverse play command.

11. An apparatus for monitoring an execution of a query, wherein the query execution comprises at least one execution thread, comprising:  
means for storing an execution trace record for each execution thread, the execution trace record having execution trace information including a thread identifier (ID) and a time stamp;  
means for retrieving the execution trace information; and  
means for presenting the retrieved execution trace information to a user after execution of the query.

12. The apparatus of claim 11, further comprising means for executing the query, and wherein means for storing an execution trace record for each execution thread is performed while executing the query.

5 13. The apparatus of claim 11, wherein the query execution comprises a plurality of execution threads, and the means for presenting the retrieved execution trace information to the user comprises means for synchronizing the execution trace records according to the thread ID and the time stamp.

10 14. The apparatus of claim 11, wherein the time stamp is an absolute time stamp.

15 15. The apparatus of claim 11, wherein the time stamp is a logical time stamp.

16. The apparatus of claim 11, wherein the means for storing an execution trace record for each execution thread in at least one execution log file comprises means for splitting the execution log file into a plurality of partitions.

20 17. The apparatus of claim 16, wherein the execution log file is split into the plurality of partitions based upon an amount of execution trace information and an amount of storage space.

25 18. The apparatus of claim 16, further comprising means for reusing the partitions when the execution log file exceeds a partition size.

19. The apparatus of claim 11, wherein the means for presenting the retrieved execution trace information to a user comprises:  
means for accepting a presentation command; and  
means for arranging the retrieved execution trace records according to the  
5 thread ID, the time stamp and the playback command.

20. The apparatus of claim 19, wherein the presentation command is selected from the group comprising a play command, a stop command, a reverse play command, a fast play command, and a fast reverse play command.

21. A program storage device, readable by a computer, tangibly embodying at least one program of instructions executable by the computer to perform method steps of monitoring an execution of a query, wherein the query execution comprises at least one execution thread, the method steps comprising the steps of:

15 storing an execution trace record for each execution thread, the execution trace record having execution trace information including a thread identifier (ID) and a time stamp;

retrieving the execution trace information; and

20 presenting the retrieved execution trace information to a user after execution of the query.

22. The program storage device of claim 21, wherein the method steps further comprise the step of executing the query, and wherein the method step of storing an execution trace record for each execution thread is performed while  
25 executing the query.

5 23. The program storage device of claim 21, wherein the query execution comprises a plurality of execution threads, and the method step of presenting the retrieved execution trace information to the user comprises the method step of synchronizing the execution trace records according to the thread ID and the time stamp.

24. The program storage device of claim 21, wherein the time stamp is an absolute time stamp.

10 25. The program storage device of claim 21, wherein the time stamp is a logical time stamp.

15 26. The program storage device of claim 21, wherein the step of storing an execution trace record for each execution thread in at least one execution log file comprises the step of splitting the execution log file into a plurality of partitions.

20 27. The program storage device of claim 26, wherein the execution log file is split into the plurality of partitions based upon an amount of execution trace information and an amount of storage space.

28. The program storage device of claim 26, further comprising the steps of reusing the partitions when the execution log file exceeds a partition size.

25 29. The program storage device of claim 21, wherein the step of presenting the retrieved execution trace information to a user comprises the steps of:  
accepting a presentation command; and  
arranging the retrieved execution trace records according to the thread ID, the time stamp and the playback command.

30. The program storage device of claim 29, wherein the presentation command is selected from the group comprising a play command, a stop command, a reverse play command, a fast play command, and a fast reverse play command.

5 31. An apparatus for monitoring an execution of a query, wherein the query execution comprises at least one execution thread, comprising:

a data server for executing the execution thread and for storing an execution trace record for the executed execution thread, the execution trace record having execution trace information including a thread identifier (ID) and a time stamp;

10 a query coordinator, for storing an execution plan having a time stamp and for retrieving and synchronizing the execution trace record and the execution plan; and

a client process for displaying the retrieved execution trace information to a user after execution of the query.

15 32. The apparatus of claim 31, wherein the data server and the query coordinator stores the execution trace record and the execution plan while executing the query.

20 33. The apparatus of claim 31, wherein the query execution comprises a plurality of execution threads, and the query coordinator synchronizes the execution trace records according to the thread ID and the time stamp.

25 34. The apparatus of claim 31, wherein the time stamp is an absolute time stamp.

35. The apparatus of claim 31, wherein the time stamp is a logical time stamp.